LESS IS NOT WORSE: EFFECTIVE REASONING WITHOUT COMPLETE REASONING TRACES

Anonymous authorsPaper under double-blind review

ABSTRACT

Large language models (LLMs) often produce lengthy reasoning traces with substantial token redundancy. While reasoning processes are widely adopted to tune LLMs as a post-training regime, it has been underexplored whether LLMs truly learn from the complete trajectory, particularly in supervised fine-tuning (SFT). We argue that, for mid-size LLMs commonly trained with SFT for reasoning, using full reasoning trajectories may harm performance because their limited capacity increases susceptibility to redundant intermediate steps. To investigate, we first analyze the redundancy in thinking trajectories through attention maps and controlled token-removal studies, both of which show that intermediate tokens contribute minimally to reasoning quality. Our analyses suggest that the most redundant segments typically appear in the middle of reasoning traces, whereas the earlier and later segments are crucial for generating high-quality final outcomes. We further posit that avoiding redundant intermediate information leads to exploiting the capability of LLMs to infer concise and coherent intermediate steps by utilizing the known start and end points. Based on the insights, we propose **MidCut**, a method that removes redundant middle steps during both training and inference. We demonstrate the effectiveness of **MidCut** in two scenarios for LLM reasoning: (1) SFT trained on s1K and OpenThoughts datasets for reasoning; and (2) decoding strategy for a test-time application.

1 Introduction

Learning to write for human learners may offer a helpful analogy for understanding how large language models (LLMs) (Grattafiori et al., 2024; Yang et al., 2024a) acquire reasoning skills. Learners benefit from carefully curated texts (*e.g.*, mentor texts), which are essential for building writing fundamentals (Kim et al., 2021; Culham, 2023; Shubitz, 2023). Yet, as learners advance, they increasingly dismiss misleading information as unhelpful, gaining little from it; they rather place more value on accurate and essential guidance, which drives more effective learning. A follow-up question would be: if provided only with essentials to learn yet disconnected parts, would they learn better (or even faster) and use prior knowledge to predict the missing pieces on their own?

LLMs have recently achieved strong results in reasoning tasks driven by large reasoning models (LRMs), such as the pioneer GPT-40 (Hurst et al., 2024) and various open-sourced initiatives (Yang et al., 2025; 2024b; Guo et al., 2025). The tasks span domains such as mathematical problem-solving and logical reasoning, which demand complex processes that are not known to be achievable solely through pre-training, mid-training, or alignment. LRMs are often trained through reward-based learning (Shao et al., 2024), or by leveraging machine-generated reasoning traces in a manner similar to how human learners are guided, to enhance reasoning capability. The reasoning traces, which represent the trajectory from a question to its answer, are informative for learning reasoning traces as supervisory signals; however, they often include redundant or unnecessary steps. Following the earlier analogy, one may ask whether a model trained only on essential sub-parts would fill the missing pieces in a trajectory with more plausible ones. However, even these are underexplored: identifying which parts of the reasoning trajectory are redundant and formulating a principle to exploit only the essential ones.

Existing studies examining the *thinking (reasoning) trajectory* of LRMs intentionally separated from their final answers to make the reasoning process explicit. Muennighoff et al. (2024) first introduced

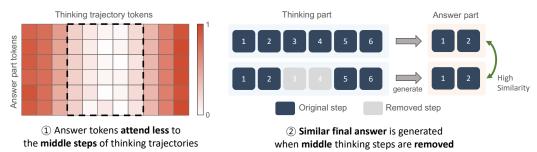


Figure 1: **Redundancy may emerge in thinking trajectories.** Our insight is that a thinking trajectory often includes less informative traces, which we term redundancy. We preview our analyses for demonstration in §3: (left) low attention weights across thinking trajectories on an intermediate set of tokens; (right) removing them yields similar outputs.

the idea of using *entire* reasoning traces as guidance for supervised fine-tuning (SFT), and subsequent datasets such as (Guha et al., 2025) have further contributed reasoning traces for learning. Subsequently, however, several studies have noted that the thinking trajectories produced by LRMs are typically lengthy, less informative, and sometimes misleading (Ma et al., 2025b; Cuesta-Ramirez et al., 2025; Wu et al., 2025). Prior work further suggests that LLMs often already know the answer before generating a fully explicit reasoning trace (Lindsey et al., 2025). Literature (Ma et al., 2025b; Wang et al., 2025) demonstrates that problems can sometimes be efficiently solved without any explicit reasoning process, albeit with lower performance sometimes. These observations suggest that reasoning trajectories are important for solving complex problems, yet the fully explicit trajectory may not always be necessary. This raises a fundamental question: *do LLMs require the complete reasoning trajectory, and if not, which parts are non-essential?*

This paper focuses on SFT for mid-size LLMs using machine-generated reasoning traces (Muennighoff et al., 2024; Team, 2025; Labs, 2025; Guha et al., 2025). Our key insight is that such traces often contain substantial redundancy, particularly in intermediate steps; removing these redundancies enables more effective learning and yields significant performance gains due to the often-limited model capacity. To further investigate our insights, we first propose two systematic analyses: attention weight patterns and knockouts motivated by prior works (Ameisen et al., 2025; Clark et al., 2019; Madaan et al., 2023; Geva et al., 2023). We find that answer tokens place little attention on middle steps, and that removing these steps preserves answer quality, in contrast to removing the beginning or end, as shown in Figure 1. These findings indicate that LLMs do not fully rely on the complete reasoning trajectory; intermediate steps are largely redundant and non-essential. Motivated by these findings, we propose MidCut, a simple approach that trims intermediate steps of reasoning trajectories. MidCut-SFT leverages trimmed trajectories for more efficient supervised fine-tuning (SFT), while MidCut-Decoding skips redundant steps at inference to accelerate reasoning generation. Experiments demonstrate that both strategies improve efficiency while consistently maintaining strong performance across diverse reasoning tasks.

2 RELATED WORK

Test-time scaling and consideration of overthinking. Muennighoff et al. (2024) introduced a small curated dataset that is dubbed s1K having 1,000 reasoning trace examples satisfying difficulty, diversity, and quality, and *budget forcing* mechanism to control reasoning length at inference (*e.g.*, forcing early stop or appending "Wait" to prolong reasoning). Their model (*e.g.*, s1-32B) exhibits strong improvements in competition math questions (*e.g.*, AIME24) and shows a clear positive scaling-trend with test-time compute. After it, some recent works have challenged the assumption that longer reasoning chains always improve performance. Ghosal et al. (2025) analyzed test-time scaling and shows that beyond a certain point, longer chains can actually *hurt accuracy* due to "overthinking" rather than enhancing reasoning. Similarly, Hassid et al. (2025) demonstrated that, when multiple chains are sampled for the same question, the *shortest chain is often more reliable than the longest*, motivating inference-time strategies such as short-m@k and fine-tuning on shorter demonstrations. However, these approaches treat *length itself* as the key factor, without considering whether different segments of a chain (beginning, intermediate, and end) contribute unequally.

 Suppressing or replacing reasoning tokens for efficiency. Another line of work explores omitting or substituting explicit 'thinking' tokens to achieve efficient reasoning. Ma et al. (2025a) found that reasoning models can remain competitive even when the entire explicit reasoning block is skipped (NoThinking). Wang et al. (2025) showed that suppressing reflection-like markers such as "Wait" or "Alternatively" significantly shortens chains while maintaining performance. This demonstrates that not all generated tokens are necessary. Still, this strategy mainly targets special filler tokens rather than the broader redundancy of standard reasoning steps within the middle of a chain. Ringel et al. (2025) proposed a learned continue-thinking token that dynamically extends reasoning when needed. Both focus on whether to engage in any reasoning at all or on how to prolong it, but do not ask whether some parts of the reasoning trajectory are inherently more redundant than others.

Reasoning compression and adaptive selectivity. A similar line of work seeks to *compress* reasoning traces or *adaptively decide* how much to reason also for efficiency. At the trajectory level, Hou et al. (2025) pruned tokens *primarily considering sequence length* via reinforcement learning with an explicit token length limit, iteratively tightening the limit to shorten thoughts with minimal accuracy loss. Fan et al. (2025) proposed that *a light instruction model* drafts a high-level outline, and a reasoning model fills in details, which reduces generated tokens while maintaining accuracy, thus enabling difficulty-aware depth adjustment. Yuan et al. (2025) introduces a token-level compression framework that *scores reasoning tokens* and trains on compacted CoT while preserving accuracy. Lin et al. (2024) demonstrated at *pretraining* time that not all tokens are equally useful: they focus loss selectively on *high-utility tokens* (*i.e.*, useful and clean tokens in their terms), improving data efficiency – conceptually aligned with compression leading to faster training. These approaches either rely on token-importance scoring, budgeted RL, or guidance policies. By contrast, our approach offers a simple *region-level* recipe at training time: systematically remove the middle span of machine-generated chains during SFT, keeping the earlier and later spans of traces while discarding the redundant middle, without the need for auxiliary scorers or controllers.

Our focus: to identify redundancy and to propose a simple solution to address it. Our work is distinguished in three ways: (1) region-level perspective: we identify and omit a specific region rather than entire chains or individual tokens; (2) training-time method: unlike test-time heuristics, our method reshapes the training data itself; (3) simple, low cost and reproducible: our approach requires no auxiliary scoring modules and scales easily across large datasets. Finally, when recast as a question, while existing research asks "How long should a trajectory be?" or "Which tokens should we selectively suppress?", this work asks "Which parts of the trajectory are essential?"

3 Analyzing thinking trajectories

This section begins by presenting some preliminaries. We then investigate whether certain tokens within reasoning trajectories are redundant. We employ two complementary approaches: an attention-based analysis of how models process reasoning traces and a knockout-based analysis of how traces affect the quality of answer generation. Analysis overview is illustrated in Figure 1.

3.1 BACKGROUND

Supervised fine-tuning (SFT)-based reasoning training. Recent methods with carefully curated data: Sky-T1 (Team, 2025), s1K (Muennighoff et al., 2024), Bespoke-Stratos (Labs, 2025), OpenThoughts (Guha et al., 2025) suggest a promising paradigm, beyond RL-based reasoning training (Shao et al., 2024): reasoning models can be effectively trained with curated, machine-generated traces produced by an LRM, for example.

Formally, let $\mathcal{D} = \{(x_i, y_i, r_i)\}_{i=1}^N$ denote a altogether data of input query and answer pairs (x_i, y_i) augmented with a reasoning trace r_i generated by a LRM T (e.g., DeepSeek-R1 671B (Guo et al., 2025)). SFT then optimizes the parameters θ of a language model M_{θ} (usually smaller than T e.g., 3B-, 7B-, or 32B- scale) by minimizing the negative log-likelihood:

$$\mathcal{L}_{SFT}(\theta) = -\sum_{i=1}^{N} \log p_{\theta}(y_i, r_i \mid x_i), \tag{1}$$

where r_i can be partially or fully included depending on the training recipe.

In this setting, high-quality reasoning traces provided by T act as rich supervision signals, enabling smaller models M_{θ} to acquire strong reasoning ability without resorting to RL-based objectives (e.g., GRPO (Shao et al., 2024)). This setup is particularly practical: large-scale traces are already well-formed by powerful teacher models, and SFT alone has been shown to yield competitive performance in smaller models.

Problem setting. We focus on SFT training with machine-generated reasoning traces (Muennighoff et al., 2024; Guha et al., 2025), applied to mid-scale models (*e.g.*, \simeq 32B parameters). While RL often surpasses SFT (Chu et al., 2025), we presume that applying SFT is reasonable in this setting, both in terms of simplicity and expected performance. Muennighoff et al. (2024)'s findings implicitly supported the assumption: when large teacher models (*e.g.*, DeepSeek-R1 671B (Guo et al., 2025)) already produce high-quality reasoning traces, it is both realistic and effective to distill these traces into smaller models through SFT. Furthermore, since *SFT typically precedes RL methods* in RLHF procedures (Ouyang et al., 2022; Liu et al., 2024; Achiam et al., 2023; Groeneveld et al., 2024) or is employed successively (OLMo et al., 2024; Yang et al., 2025), understanding the mechanics of SFT in relation to reasoning can give valuable insights. Therefore, we believe it would be valuable to study *how reasoning traces can be more effectively exploited for SFT training*.

Our insight is threefold: (1) redundancy (within reasoning chains of thinking trajectories) makes full traces difficult to learn with FT, particularly for models with limited capacity; (2) intermediate steps in reasoning traces are often redundant; and (3) removing them allows the model to focus on filling them with more relevant content to reach the final answer. Grounded in our insights, we conjecture that eq. (1) would become increased when some parts of reasoning trajectories are less informative or redundant. Before studying how redundancy affects the empirical behavior of LLMs, we first analyze redundancy in thinking trajectories.

3.2 Analyses

Figure 2 illustrates both example instances and the overall structure of reasoning trajectories. Such trajectories typically proceed from problem definition, through exploratory reasoning, to final answer consolidation. To investigate the functional roles of different parts in reasoning trajectories, we examine representative examples across reasoning tasks. Intriguingly, as similarly noted in prior work (Guo et al., 2025; Ma et al., 2025b; Cuesta-Ramirez et al., 2025; Wu et al., 2025), we often observe redundant traces (*e.g.*, repeated checks, backtracking, or unnecessary elaborations) in the middle of a trajectory. We now systematically investigate whether such human-perceived redundancy indeed matters for LLM.

Attention weights analysis. Following prior works (Ameisen et al., 2025; Clark et al., 2019; Madaan et al., 2023; Geva et al., 2023) that interpreted models using attention-based metrics, we analyze attention weights to investigate how different parts of thinking trajectories contribute to answer generation. We expect these patterns to reveal how models prioritize tokens during generation, offering insights into information flow within transformer architectures.

Figure 3 provides two complementary insights. First, the overall (*i.e.*, averaged across layers) trend in Figure 3(a) highlights **strong attention peaks at the beginning and ending of trajectories**, whereas intermediate steps are weakly attended. This suggests that intermediate steps generally contribute less to answer generation, indicating that these tokens could be redundant within the overall reasoning trajectory. Second, the layer-specific patterns in Figure 3(b–e) reveal a progression across the model depth: very early layers (pattern 1) attend broadly without clear localization, early layers (pattern 2) shift focus toward candidate answer tokens, intermediate layers (pattern 3) balance attention between early problem-definition and later reasoning steps, and final layers (pattern 4)

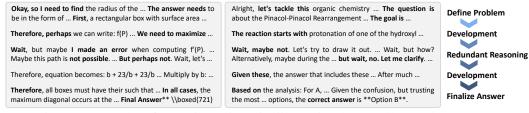


Figure 2: Examples of reasoning trajectories illustrating the overall reasoning process.

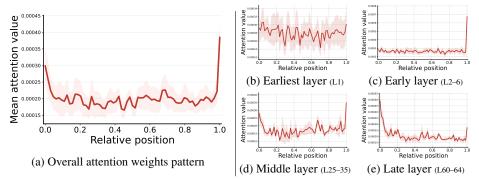


Figure 3: **Averaged 1D attention weights across reasoning trajectories** during answer generation. (a) The overall average highlights strong attention peaks at the beginning and end of trajectories. (b–e) Layer-specific patterns reveal a progression: very early layers attend broadly, early layers shift their focus toward candidate answers, intermediate layers jointly reference problem-definition and reasoning steps, and final layers concentrate on the beginning to format and consolidate the answer.

concentrate strongly on the beginning to format and consolidate the final answer. Taken together, these findings suggest how information flow evolves through the network and provide evidence that intermediate reasoning steps are deemphasized as generation progresses toward the later layer.

Attention knockout analysis. We further apply the attention knockout technique (Geva et al., 2023) for our analysis, which masks attention links to probe the importance of specific trajectory segments. Specifically, we truncate the beginning (e.g., the first 0–10% of tokens), the intermediate (e.g., a centered span such as 45–55%), or the ending (e.g., the last 90–100%) and compare the resulting answers against those from the full trajectories. This allows us to retest which part of reasoning trajectories are causally important for answer generation. Additional details on the experimental setup and evaluation metric are provided in Appendix A.1.

Based upon the observational insights from attention analysis, Figure 4 shows the average results over samples across different knockout ratios. In the results, **removing the intermediates consistently yields the highest similarity** to answers generated with the full trajectory. In contrast, removing the beginning and ending parts results in lower similarity, suggesting that these segments contain critical information for formulating correct answers. This provides quantitative evidence for our hypothesis that intermediate steps are often redundant and do not substantially contribute to final answer quality.

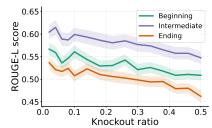


Figure 4: **ROUGE-L under attention knockouts** for each segment of reasoning traces. See more in Figure A.1.

Takeaway from §3. ■ *Beginning and ending segments matter most, intermediate less:* Both attention and knockout analyses show that LLMs assign less importance to intermediate parts.
2 *Answer-first, then question-refinement:* attention patterns reveal an information flow in which models first lean on the final answer and then revisit the question for refinement.



Figure 5: Overview of MidCut-SFT and MidCut-Decoding methodology. MidCut-SFT processes training data by removing middle parts from thinking trajectories, while MidCut-Decoding applies the same principle during inference to reduce computational overhead.

4 TRIMMING REASONING TRACES TO ENHANCE TRAINING AND DECODING

Based on the knowledge in §3, which reveals redundancy in the middle parts of thinking trajectories, we introduce <code>MidCut</code>, a method for removing intermediate steps from reasoning trajectories. As illustrated in Figure 5, we present two applications of our <code>MidCut</code> methodology: <code>MidCut-SFT</code> is a preprocessing method applied to LRM-generated reasoning trajectories in SFT datasets. This approach removes intermediate parts of the trajectory while preserving the beginning and ending segments, which are believed to contain the essential setup of the problem and the conclusion. <code>MidCut-Decoding</code> applies the same principle during inference, truncating intermediate reasoning steps after the thinking phase to reduce computational overhead and generation latency without compromising the quality of final answers. In both applications, our method targets the redundant middle parts identified in our analysis while preserving critical components at the trajectory boundaries.

4.1 MIDCUT-SFT: MIDCUT FOR SUPERVISED FINE-TUNING

The proposed method. MidCut-SFT follows a simple yet effective approach: given reasoning trajectories for training, we remove the intermediate traces. Specifically, we preserve the early and late segments of the reasoning trajectories based on predefined thresholds, either by step count or token count. Steps are defined by splitting trajectories at double newline characters $(\n\n)$, represented by natural reasoning breakpoints.

We explore several variants of this filtering approach: (1) Step-level filtering: Remove intermediate steps while preserving the first and last n steps of the trajectory. (2) Token-level filtering: Remove intermediate tokens while preserving the first and last k tokens, regardless of step boundaries. (3) Length-proportional filtering: We remove a percentage of steps proportional to the original trajectory length, specifically removing m% of the total steps from the middle part. This approach scales the removal amount based on trajectory complexity. (4) Similarity-based filtering: We compute Jaccard similarity between each step and the preceding 5 steps of the trajectory. Steps exceeding a predefined similarity threshold are filtered out to reduce redundancy, targeting repetitive reasoning patterns. Further details on the truncation segments and parameter settings are provided in Appendix B.3.

Comparison methods. To evaluate the effectiveness of our approach, we compare it against several alternative trajectory reduction strategies: (1) LLM-based Compression: we use external LLMs –

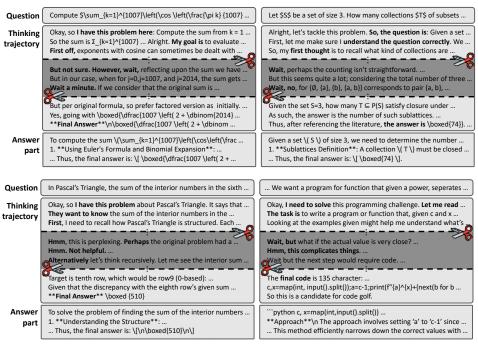


Figure 6: Examples of MidCut-SFT applied to (top) s1K-1.1 and (bottom) OpenThoughts3datasets.

Table 1: **Main comparison results** of **MidCut**-SFT on AIME24, GPQA-D, and MATH across different models and datasets. Positive relative changes are shown in blue, negative in red, and the best results for each task are highlighted in bold.

M-41 1	AII	ME24	GP	QA-D	MATH		Average				
Method	Value	Δ	Value	Δ	Value	Δ	Value	Δ			
Qwen2.5-32B-Instruct & s1K-1.1											
Base	0.6444	_	0.6195	_	0.9413	_	0.7351	_			
LLM-based	0.3333	-48.28%	0.5791	-6.52%	0.8900	-5.45%	0.6008	-18.25%			
Random	0.5667	-12.07%	0.6162	-0.54%	0.9447	+0.35%	0.7092	-3.53%			
Ours	0.6889	+6.90%	0.6229	+0.54%	0.9440	+0.28%	0.7519	+2.29%			
Qwen3-8B-Base & s1K-1.1											
Base	0.4222	_	0.5741	_	0.9210	_	0.6391	_			
LLM-based	0.1556	-75.85%	0.4747	-23.37%	0.8233	-12.52%	0.4845	-34.10%			
Random	0.2833	-32.89%	0.5741	+0.00%	0.8885	-3.53%	0.5820	-8.94%			
Ours	0.4444	+5.26%	0.5741	+0.00%	0.9160	-0.54%	0.6448	+0.90%			
	Qwen3-8B-Base & OpenThoughts3-100K										
Base	0.5000	_	0.5387	_	0.9347	_	0.6578	_			
Random	0.4667	-6.67%	0.5312	-1.40%	0.9254	-1.00%	0.6411	-2.54%			
Ours	0.5222	+4.44%	0.5640	+4.70%	0.9387	+0.43%	0.6750	+2.61%			
Qwen3-4B-Base & OpenThoughts3-100K											
Base	0.3667	_	0.4865	_	0.9160	_	0.5897				
Random	0.3444	-6.08%	0.4680	-3.80%	0.9127	-0.36%	0.5750	-2.49%			
Ours	0.4000	+9.08%	0.4899	+0.70%	0.9113	-0.51%	0.6004	+1.81%			

Claude Sonnet¹ (Anthropic, 2025) and Gemini (gemini-2.5-flash) (Comanici et al., 2025) – to compress thinking trajectories while preserving essential reasoning content. In the main text we report results with Claude, while detailed results with both Claude and Gemini are provided in Appendix B.4. (2) Random Step Selection: we randomly sample 2n steps from the original trajectory, preserving the same amount of content but without structural consideration. (3) Single-end Preservation: instead of preserving both beginning and end parts, we retain either the first 2n steps or the last 2n steps, maintaining the same total preserved length as **MidCut-SFT**.

Experimental settings. We conduct experiments on two reasoning SFT datasets: s1K-1.1 (Muennighoff et al., 2024) and OpenThoughts3-1.2M (Guha et al., 2025). For OpenThoughts3, we subsample 100K examples from the full 1.2M corpus, and further restrict to instances where both the begin-of-think (<think>) and end-of-think (</think>) tokens are present. On s1K-1.1, we fine-tune Qwen2.5-32B-Instruct and Qwen3-8B-Base for 5 epochs, while on OpenThoughts3 we fine-tune Qwen3-8B-Base and Qwen3-4B-Base for 1 epoch. More detailed settings are described in Appendix B.3. For evaluation, we use the AIME24, GPQA-D, and MATH benchmarks.

Experimental results. Table 1 shows the overall results that **MidCut**-SFT improves accuracy while using substantially fewer training tokens, with trimming the middle of thinking trajectories even outperforming training on full trajectories. Specifically, training on s1K-1.1 uses 21.9% fewer tokens and training on OpenThoughts3 uses 18.1% fewer tokens compared to full-trajectory SFT.

Table 2 shows the extended comparison results on s1K-1.1 against alternative reductions, and further detailed results are provided in Appendix B.4. LLM-based compression performs noticeably worse because it disrupts patterns of word usage and reasoning structure. Furthermore, it introduces nonnegligible computational costs and lacks precise control over the compression ratio. In contrast, approaches that preserve the original text but selectively trim segments generally achieve better performance than LLM-based compression. Among them, keeping both the beginning and ending while removing the middle consistently outperforms single-end truncation or random selection. This confirms that the location of truncation matters. Overall, these results support our hypothesis that intermediate steps are largely redundant and can even hinder effective learning if retained.

Table 3 reports ablations of **MidCut**-SFT. Step-level trimming, which preserves the beginning and ending segments at the semantic level, consistently outperforms token-level trimming. Length-

¹We utilize the claude-sonnet-4-20250514 version of Claude Sonnet 4.

379

380

381 382

389

396

397

398

399 400

401 402

403

404

405

406

407

408

409

410

411

412

413

414

415

416

417 418 419

420 421

422

423

424

425

426

428

429

430

431

Table 2: Single-end preservation vs. MidCut-SFT (Step-level) and other strategies in Table 1 on s1K-1.1. Each value represents the average performance on AIME24, GPQA-Diamond, and MATH. Relative changes from the base are shown in parentheses below each value.

Model	Base	LLM-based	Random	Ours	Single (first)	Single (last)
Qwen2.5-32B	0.7351	0.6008 (-18.25%)	0.7092 (-3.53%)	0.7519 (+2.29%)	0.7367 (+0.22%)	0.7183 (-2.30%)
Qwen3-8B	0.6391	0.4845 (-34.10%)	0.5820 (-8.94%)	0.6448 (+0.90%)	0.6099 (-4.56%)	0.5698 (-10.82%)

Table 3: Comparison among trimming variants of MidCut-SFT on s1K-1.1. Each value represents the average performance on AIME24, GPQA-Diamond, and MATH. Relative changes from the base are shown in parentheses below each value. "Len-prop." denotes length-proportional filtering.

Model	Base	Step-level	Token-level	Len-prop.	Similarity
Qwen2.5-32B	0.7351	0.7519 (+2.29%)	0.7264 (-1.19%)	0.7470 (+1.62%)	0.7392 (+0.56%)
Qwen3-8B	0.6391	0.6448 (+0.90%)	0.6162 (-3.59%)	0.6040 (-5.49%)	0.5994 (-6.21%)

proportional or similarity-based variants require additional complexity, but they do not surpass the simpler step-level MidCut. These results indicate that step-level MidCut-SFT is both simple and effective, making it preferable to more complex alternatives.

4.2 MIDCUT-DECODING: MIDCUT FOR EFFECTIVE DECODING

Method. MidCut-Decoding applies the middle-cutting principle during inference to reduce computational overhead. After the model completes its thinking trajectory, we cut the middle part of the trajectory before producing the final answer. We implement four ra-Table 4: **Performance of**

tios with {25, 33, 50, 75}% of the middle part while preserving beginning and ending segments. More details are reported in Appendix C.

MidCut-Decoding. Method Accuracy Full trajectory 0.6578 MidCut 25% 0.6574 MidCut 33% 0.6579 MidCut 50% 0.6581 MidCut 75%

0.6576

Experimental results. Table 4 reports **MidCut**-Decoding results on the SFT-trained model with OpenThought3. Compared to using the full reasoning trajectory, trimming 25–75% of the intermediate steps yields almost identical average performance. These findings indicate that MidCut-Decoding can substantially reduce inference costs for long reasoning traces while maintaining comparable task accuracy.

Takeaway from §4. Less is not worse, sometimes even better.

Removing redundant intermediate steps with **MidCut**-SFT improves reasoning performance, while inherently reducing training cost. Likewise, MidCut-Decoding sustains reasoning quality while improving efficiency in answer generation by removing redundant middle parts.

5 **DISCUSSION**

This section initiates discussions by posing crucial questions that help us gain a deeper understanding of our work. We first validate our claim that training loss is indeed reduced, and subsequently discuss what our method achieves indeed and under which conditions it is effective.

Does MidCut-SFT lead to a better convergence point? Here, we evaluate whether our conjecture holds by examining the loss curves during training on OpenThoughts3, with and without **MidCut**-SFT. As shown in Figure 7, our method consistently achieves lower perplexity across training while using fewer tokens. This finding supports our claim in §3 that redundant intermediate steps not only prolong trajectories but also inflate the training loss, which presumably makes optimization harder for LLMs; thus, LLMs would converge to a suboptimal

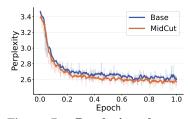


Figure 7: **Perplexity changes** when applying MidCut-SFT.

Figure 8: Token length changes.

433

434

435

436

437

438

439

440 441

442

443 444

445

446

448

449

450

451

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468

469

470

471

472

473 474

475 476

477

478

479

480

482

483

484

485

Figure 9: Impact of MidCut-SFT on intermediate traces.

point. By trimming such redundancy, the model can allocate its capacity to more informative reasoning signals, which in turn facilitates more efficient and stable learning dynamics.

What does MidCut-SFT do? Building on the loss analysis, we next examine how MidCut-SFT alters the actual reasoning behavior of models. As shown in Figure 8, we observe a consistent pattern: the number of reasoning tokens with MidCut-SFT adjusts depending on the baseline trajectories without MidCut-SFT. When original trajectories are excessively long (>15k tokens), the output length decreases; for relatively shorter ones, token usage increases. This adaptive adjustment is further illustrated with the examples in Figure 9. In the upper row, where the baseline model produces unnecessarily lengthy and redundant output, MidCut-SFT suppresses such detours, enabling the model to reach a more accurate solution with fewer tokens. In the lower row, although both models arrive at the same intermediate equations, the baseline fails to solve the problem, whereas MidCut-SFT succeeds by more effectively completing the missing steps in the reasoning process. These observations suggest MidCut-SFT enhances performance by (1) reducing redundant reasoning patterns and (2) improving the model's ability to carry out intermediate reasoning when needed.

When does MidCut-SFT work? Figure 10 plots the relationship between general performance (MMLU-Pro and MMLU-redux) without any additional training and performance gain from applying MidCut-SFT. Each point corresponds to a different base model, with results shown for both s1K-1.1 and OpenThoughts3. The detailed experimental settings are provided in Appendix B.5. We find that the effectiveness of MidCut-SFT depends strongly on the original capability of the base model. For s1K-1.1, improvements appear only when the base performance exceeds roughly 60–65, whereas the larger and more diverse OpenThoughts3 detaset shifts this threshold

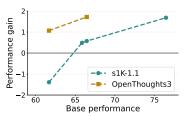


Figure 10: Performance gain by **MidCut**-SFT.

and more diverse OpenThoughts3 dataset shifts this threshold lower, allowing benefits to emerge even for weaker base models. This trend aligns with our earlier discussion: MidCut-SFT enhances the model's ability to complete intermediate reasoning steps, but this requires a sufficient baseline capacity. When models are too weak, they lack the ability to reliably infer the trimmed parts, and the method brings little to no benefit. By contrast, once a model has reached a moderate competence level, trimming redundant reasoning forces it to allocate capacity toward informative steps, thereby yielding measurable gains. This pattern closely parallels our analogy to human learners: while novices depend on full guidance, more advanced learners can benefit from streamlined instruction and effectively interpolate the missing reasoning by leveraging prior knowledge.

6 Conclusion

We have studied the necessity of complete reasoning trajectories and proposed MidCut for trajectory reduction. Through our analyses, we have shown that intermediate reasoning traces are often redundant. MidCut-SFT achieved the improved reasoning results further with efficiency, and MidCut-Decoding reduced inference costs without compromising quality. Beyond SFT and decoding, our method may also be integrated into reinforcement learning (RL) approaches such as GRPO, where reducing redundant reasoning trajectories could further enhance reasoning performance. We believe that MidCut could serve as an effective initialization for RL-based methods, and that the insights from our study may provide complementary guidance for their development. Overall, our simple yet effective method, which showcases how to handle lengthy reasoning trajectories without relying on any challenging restrictions, provides a promising research direction and paves the way for developing more efficient reasoning models in the near future.

Reproducibility statement. We performed all experiments using the s1K-1.1 and OpenThoughts3 datasets and publicly available LLMs from the Qwen series. Detailed hyperparameters and training schedules are included in the Appendix, and all evaluations were repeated across multiple random seeds, with averaged results reported. These efforts ensure that our results can be readily reproduced.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023. 4
- Emmanuel Ameisen, Jack Lindsey, Adam Pearce, Wes Gurnee, Nicholas L. Turner, Brian Chen, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. Circuit tracing: Revealing computational graphs in language models. *Transformer Circuits Thread*, 2025. URL https://transformer-circuits.pub/2025/attribution-graphs/methods.html. 2, 4
- Anthropic. Claude 4. https://www.anthropic.com/claude/sonnet, 2025. 7, 14
- Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation model post-training. *arXiv preprint arXiv:2501.17161*, 2025. 4
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. What does bert look at? an analysis of bert's attention. *arXiv preprint arXiv:1906.04341*, 2019. 2, 4
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv* preprint arXiv:2507.06261, 2025. 7, 14
- Jhouben Cuesta-Ramirez, Samuel Beaussant, and Mehdi Mounsif. Large reasoning models are not thinking straight: on the unreliability of thinking trajectories. *arXiv preprint arXiv:2507.00711*, 2025. 2, 4
- Ruth Culham. Writing thief: Using mentor texts to teach the craft of writing. Routledge, 2023. 1
- Siqi Fan, Peng Han, Shuo Shang, Yequan Wang, and Aixin Sun. Cothink: Token-efficient reasoning via instruct models guiding reasoning models. *arXiv preprint arXiv:2505.22017*, 2025. 3
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. Dissecting recall of factual associations in auto-regressive language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. 2, 4, 5
- Soumya Suvra Ghosal, Souradip Chakraborty, Avinash Reddy, Yifu Lu, Mengdi Wang, Dinesh Manocha, Furong Huang, Mohammad Ghavamzadeh, and Amrit Singh Bedi. Does thinking more always help? understanding test-time scaling in reasoning models. *arXiv preprint arXiv:2506.04210*, 2025. 2
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 1
- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, et al. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*, 2024. 4
- Etash Guha, Ryan Marten, Sedrick Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, et al. Openthoughts: Data recipes for reasoning models. *arXiv preprint arXiv:2506.04178*, 2025. 2, 3, 4, 7

- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv* preprint arXiv:2501.12948, 2025. 1, 3, 4
 - Michael Hassid, Gabriel Synnaeve, Yossi Adi, and Roy Schwartz. Don't overthink it. preferring shorter thinking chains for improved llm reasoning. *arXiv preprint arXiv:2505.17813*, 2025. 2
 - Bairu Hou, Yang Zhang, Jiabao Ji, Yujian Liu, Kaizhi Qian, Jacob Andreas, and Shiyu Chang. Thinkprune: Pruning long chain-of-thought of llms via reinforcement learning. *arXiv preprint arXiv:2504.01296*, 2025. 3
 - Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. arXiv preprint arXiv:2410.21276, 2024.
 - Young-Suk Grace Kim, Dandan Yang, Marcela Reyes, and Carol Connor. Writing instruction improves students' writing skills differentially depending on focal instruction and children: A meta-analysis for primary grade students. *Educational Research Review*, 34:100408, 2021. 1
 - Bespoke Labs. Bespoke-stratos: The unreasonable effectiveness of reasoning distillation. www.bespokelabs.ai/blog/bespoke-stratos-the-unreasonable-effectiveness-of-reasoning-distillation, 2025. Accessed: 2025-01-22. 2, 3
 - Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, 2004. 13
 - Zhenghao Lin, Zhibin Gou, Yeyun Gong, Xiao Liu, Ruochen Xu, Chen Lin, Yujiu Yang, Jian Jiao, Nan Duan, and Weizhu Chen. Not all tokens are what you need for pretraining. In *NeurIPS*, 2024.
 - Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. On the biology of a large language model. *Transformer Circuits Thread*, 2025. URL https://transformer-circuits.pub/2025/attribution-graphs/biology.html. 2
 - Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv* preprint *arXiv*:2412.19437, 2024. 4
 - Wenjie Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. Reasoning models can be effective without thinking. *arXiv preprint arXiv:2504.09858*, 2025a. 3
 - Wenjie Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. Reasoning models can be effective without thinking. *arXiv preprint arXiv:2504.09858*, 2025b. 2, 4
 - Aman Madaan, Katherine Hermann, and Amir Yazdanbakhsh. What makes chain-of-thought prompting effective? a counterfactual study. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 1448–1535, 2023. 2, 4
 - Christopher D Manning. Introduction to information retrieval. Syngress Publishing,, 2008. 13
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candes, and Tatsunori Hashimoto. s1: Simple test-time scaling. In *Workshop on Reasoning and Planning for Large Language Models*, 2024. 1, 2, 3, 4, 7, 13
 - Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, et al. 2 olmo 2 furious. *arXiv preprint arXiv:2501.00656*, 2024. 4

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:
 27730–27744, 2022. 4

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 311–318, 2002. 13

- Liran Ringel, Elad Tolochinsky, and Yaniv Romano. Learning a continue-thinking token for enhanced test-time scaling. arXiv preprint arXiv:2506.11274, 2025. 3
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024. 1, 3, 4
- Stacey Shubitz. Craft moves: Lesson sets for teaching writing with mentor texts. Routledge, 2023.
- NovaSky Team. Sky-t1: Train your own o1 preview model within \$450. https://novasky-ai.github.io/posts/sky-t1, 2025. Accessed: 2025-01-09. 2, 3
- Chenlong Wang, Yuanning Feng, Dongping Chen, Zhaoyang Chu, Ranjay Krishna, and Tianyi Zhou. Wait, we don't need to" wait"! removing thinking tokens improves reasoning efficiency. arXiv preprint arXiv:2506.08343, 2025. 2, 3
- Yuyang Wu, Yifei Wang, Ziyu Ye, Tianqi Du, Stefanie Jegelka, and Yisen Wang. When more is less: Understanding chain-of-thought length in llms. *arXiv preprint arXiv:2502.07266*, 2025. 2, 4
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. Qwen2 technical report. arXiv preprint arXiv:2407.10671, 2024a. 1
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024b. 1
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. arXiv preprint arXiv:2505.09388, 2025. 1, 4
- Hang Yuan, Bin Yu, Haotian Li, Shijun Yang, Christina Dan Wang, Zhou Yu, Xueyin Xu, Weizhen Qi, and Kai Chen. Not all tokens are what you need in thinking. *arXiv preprint arXiv:2505.17827*, 2025. 3

Appendix

This appendix provides additional experimental details and results that complement the main paper. In Appendix A, we elaborate on the experimental settings and extended experimental results of our attention knockout analyses. Appendix B presents further implementation details and additional results of MidCut-SFT, including comparisons on the OpenThoughts3-100K dataset and results with alternative LLM-based compression methods. We also report extended results of Table 1 that complement the tables in the main text. Appendix C provides implementation details of MidCut-Decoding, specifying how the middle segments are truncated during inference. Finally, Appendix D clarifies the limited role of external LLMs in this work, restricted solely to writing polish and coding assistance, with no involvement in research ideation, implementation, or analysis.

A DETAILS OF ATTENTION KNOCKOUT ANALYSIS

A.1 EXPERIMENTAL SETTINGS

We delete specific trajectory segments and generate answers under these conditions to verify again whether thinking trajectories contain redundancy with respect to answer generation. By observing how answer quality changes under such interventions, we can assess whether particular segments play a causal role in generating answers. Our analysis is conducted on 100 problems from the GPQA-D dataset using the S1.1-32B (Muennighoff et al., 2024) model². For each problem, a full thinking trajectory is first generated. Then, we remove three different segments: the beginning segment, an intermediate segment centered in the middle, and the ending segment. To systematically examine the effect of segment length, we vary the truncated ratio (e.g., 10%, 20%). For instance, with a 10% ratio the beginning corresponds to tokens 0–10%, the intermediate to 45–55%, and the ending to 90–100% of the trajectory. With a 20% ratio, the corresponding spans become 0–20%, 40–60%, and 80–100%, respectively. The model is subsequently prompted to generate answers from these truncated trajectories as well as from the full trajectory. Finally, textual similarity between answers from truncated and full trajectories is measured using Jaccard similarity (Manning, 2008) and ROUGE-L (Lin, 2004).

A.2 ADDITIONAL RESULTS

Complementing the ROUGE-L (Lin, 2004) results in Figure 4 of the main paper, Figure A.1 reports Jaccard similarity (Manning, 2008) and BLEU score (Papineni et al., 2002) under the same knockout settings. The trends are closely similar to those of the ROUGE-L score. Across all knockout ratios (0–50%), removing the intermediate segment consistently yields the highest similarity between answer parts generated from truncated trajectories and their counterparts generated from the full trajectory. In contrast, removing the ending segment yields the lowest similarity and is even worse than removing the beginning. This aligns with the information flow revealed by the attention weight analysis (Figure 3). Models appear to anchor the answer first and then revisit the question for refinement, so truncating the final consolidation region most harms the generated answer parts.

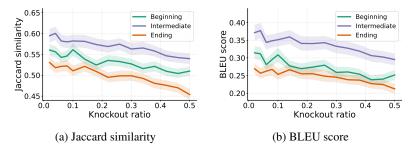


Figure A.1: (a) **Jaccard similarity and (b) BLEU score** of generated answer parts after knockouts of thinking trajectories. They also show similar trends to Figure 4.

²https://huggingface.co/simplescaling/s1.1-32B

B DETAILS OF MIDCUT-SFT EXPERIMENTS

B.1 SFT AND EVALUATION FRAMEWORKS

For the **s1K-1.1 experiments**, we conduct supervised fine-tuning using the s1 GitHub repository³. Evaluation relies on the built-in implementation of lm-evaluation-harness⁴ within the same repository. The dataset is $s1K-1.1^5$.

For the **OpenThought3 experiments**, we follow the instructions of the OpenThought GitHub repository⁶ and train models using LLaMA-Factory⁷. Evaluation is performed with lm-evaluation-harness². The dataset is based on OpenThoughts3-1.2M⁸, from which we randomly sample 100K examples containing both <think> and </think> tokens.

B.2 LLM-BASED COMPRESSION METHOD

We also experiment with compressing the s1K-1.1 dataset using two external LLMs: Claude Sonnet (claude-sonnet-4-20250514 version of Claude Sonnet 4) (Anthropic, 2025) and Gemini (gemini-2.5-flash) (Comanici et al., 2025).

To ensure that only the middle portion is shortened while respecting its local context, we operate on the thinking trajectory with a block-wise, context-aware paraphrasing pipeline. Each trajectory is split into "blocks" by "\n\n", and we iterate over non-overlapping windows of 10 consecutive blocks. For every window, we construct a prompt that exposes the three preceding and three following blocks as [Context Before] and [Context After], while designating the 10-block window as [Center Section]. The LLM is instructed to compress the [Center Section] only, preserving the logical structure and final conclusion. We instruct the LLMs with a fixed paraphrasing prompt, which explicitly conditions on surrounding context while compressing only the center section. The full prompt is shown in the box below.

Paraphrasing prompt

You will be given a step-by-step reasoning process written by a large language model.

Please paraphrase the **center section only**, while preserving the logical structure.

Try to reduce its length to approximately half of the original, but make sure to keep all essential reasoning steps and the final conclusion.

You may refer to the surrounding context for understanding, but do not modify them.

```
[Context Before]
{before}

[Center Section]
{center}

[Context After]
{after}
```

Paraphrased (center section only):

https://github.com/simplescaling/s1

⁴https://github.com/EleutherAI/lm-evaluation-harness

⁵https://huggingface.co/datasets/simplescaling/s1K-1.1_tokenized

⁶https://github.com/open-thoughts/open-thoughts

⁷https://github.com/hiyouga/LLaMA-Factory

 $^{^{8} \}verb|https://huggingface.co/datasets/open-thoughts/OpenThoughts3-1.2M|$

B.3 IMPLEMENTATION DETAILS AND EXPERIMENTAL SETTINGS

In all variants of **MidCut**-SFT, the middle portion of the trajectory is removed after preserving both the beginning and ending segments. Concretely, for step-level filtering we preserve the first and last n steps and remove the middle (total steps -2n) steps; for token-level filtering we preserve the first and last k tokens and remove the middle (total tokens -2k) tokens; and for length-proportional filtering we remove m% of the total steps from the middle part. If the preserved length (2n steps or 2k tokens) exceeds the total trajectory length, no truncation is applied.

For step-level filtering, we set $n{=}100$ for s1K-1.1 (average 234 steps) and $n{=}200$ for OpenThought3 (average 461 steps), reflecting that the latter trajectories are roughly twice as long. Step-level midcut is used as the default configuration, while further exploratory ablations (token-level with k=8000 and length-proportional with m=0.15) are carried out on s1K-1.1. For similarity-based ${\bf MidCut}$, we experimented with different filtering criteria (max, min, and average Jaccard similarity against a threshold) and found that the best performance was obtained with an average threshold of 0.3 for Qwen3-8B-Base and a minimum threshold of 0.15 for Qwen2.5-32B-Instruct, which we report in the main results. For evaluation, we use AIME24, GPQA-D, and MATH benchmarks, generating with temperature 0.6. Each instance is evaluated three times, and we report the averaged results.

B.4 ADDITIONAL RESULTS OF MIDCUT-SFT

Table B.1 compares the performance of **MidCut**-SFT on OpenThoughts3-100K against alternative reduction strategies. Consistent with the results on s1K-1.1 (Table 2), the step-level **MidCut**-SFT achieves the highest performance among all methods, yielding relative improvements of +2.61% on Qwen3-8B-Base and +1.81% on Qwen3-4B-Base compared to training on full trajectories.

Table B.2 presents detailed results by task, model, and method, complementing Table 1-3 and Table B.1. Overall, the results confirm that **MidCut**-SFT consistently achieves strong performance.

Table B.1: Comparison between **MidCut**-SFT (Step-level) and alternative truncation strategies on OpenThoughts3-100K. Each value represents the average performance on AIME24, GPQA-Diamond, and MATH. Relative changes from the base are shown in parentheses below each value.

Model	Base	Random	Ours (Step-level)	Single (first)	Single (last)
Qwen3-8B	0.6578	0.6411 (-2.54%)	0.6750 (+2.61%)	0.6216 (-5.50%)	0.5986 (-9.00%)
Qwen3-4B	0.5897	0.5750 (-2.49%)	0.6004 (+1.81%)	0.5816 (-1.38%)	0.5268 (-10.68%)

B.5 EXPERIMENTAL DETAILS OF FIGURE 10

To examine under which conditions MidCut-SFT is most effective, we analyze the correlation between the general task performance of the pretrained models and the performance gain of MidCut-SFT over full-trajectory SFT. For this experiment, we train Qwen2.5-32B-Instruct, Qwen3-8B-Base, and Qwen3-4B-Base on s1K-1.1, and Qwen3-8B-Base and Qwen3-4B-Base on OpenThoughts3-100K. The results are then compared to investigate how model capacity and initialization relate to the observed performance gains.

C IMPLEMENTATION DETAILS OF MIDCUT-DECODING

The model first generates the entire thinking trajectory up to the point where it would normally begin producing the final answer. Before the answer part is generated, we truncate the central portion of the thinking trajectory according to a specified cutting ratio. Formally, given a trajectory of length L tokens, we remove $\lfloor r \cdot L \rfloor$ tokens centered at position L/2, thereby preserving both the first (1-r)/2 and last (1-r)/2 fractions of the trajectory. For example, with r=0.5, the central 50% of tokens (i.e., from 25% to 75% of the trajectory) is removed. This ensures that the model generates its final

Table B.2: **Detailed performance of MidCut-SFT on three benchmarks (AIME24, GPQA-D, MATH).** Bold indicates the best performance in each benchmark, and underline indicates performance better than baseline. "Len-prop." denotes length-proportional filtering.

Method	AIN	ME24	GPO	QA-D	M	MATH		Average		
	Value	Δ	Value	Δ	Value	Δ	Value	Δ		
Qwen2.5-32B-Instruct & s1K-1.1										
Base	0.6444	_	0.6195	_	0.9413	-	0.7351	_		
LLM-based (Claude)	0.3333	-48.27%	0.5791	-6.52%	0.8900	-5.45%	0.6008	-18.27%		
LLM-based (Gemini)	0.1778	-72.41%	0.4815	-22.27%	0.8273	-12.12%	0.4955	-32.59%		
Single (first)	0.6222	-3.44%	0.6380	+2.99%	0.9500	+0.92%	0.7367	+0.22%		
Single (last)	0.6000	-6.90%	0.6128	-1.08%	0.9420	+0.07%	0.7183	-2.29%		
Random	0.5667	-12.06%	0.6162	-0.53%	<u>0.9447</u>	+0.36%	0.7092	-3.52%		
Ours										
Step-level	0.6889	+6.90%	0.6229	+0.55%	0.9440	+0.29%	0.7519	+2.29%		
Token-level	0.6111	-5.16%	0.6195	0.00%	0.9487	+0.79%	0.7264	-1.17%		
Len-prop.	0.6556	+1.74%	0.6380	+2.99%	$\overline{0.9473}$	+0.64%	0.7470	+1.62%		
Sim-based	0.6444	0.00%	0.6279	+1.36%	0.9453	+0.42%	0.7392	+0.56%		
	Qwen3-8B-Base & s1K-1.1									
Base	0.4222	_	0.5741	_	0.9210	_	0.6391	_		
LLM-based (Claude)	0.1556	-75.85%	0.4747	-23.37%	0.8233	-12.52%	0.4845	-34.10%		
LLM-based (Gemini)	0.1111	-73.68%	0.4293	-25.22%	0.7813	-15.16%	0.4406	-31.06%		
First	0.3667	-13.13%	0.5640	-1.76%	0.8990	-2.39%	0.6099	-4.55%		
Last	0.2778	-34.23%	0.5387	-6.17%	0.8930	-3.04%	0.5698	-10.80%		
Random	0.2833	-32.94%	0.5741	+0.00%	0.8885	-3.53%	0.5820	-8.94%		
Ours										
Step-level	0.4444	+5.26%	<u>0.5741</u>	+0.00%	0.9160	-0.54%	0.6448	+0.90%		
Token-level	0.3667	-13.13%	0.5690	-0.89%	0.9130	-0.87%	0.6162	-3.58%		
Len-prop.	0.3333	-21.00%	0.5606	-2.35%	0.9180	-0.33%	0.6040	-5.49%		
Sim-based	0.3444	-18.42%	0.5539	-3.52%	0.9000	-2.28%	0.5994	-6.22%		
		Qwen3-8l	B-Base &	OpenThou	ght3-100K	-				
Base	0.5000	_	0.5387	_	0.9347	_	0.6578	_		
First	0.4000	-20.00%	0.5320	-1.24%	0.9327	-0.21%	0.6216	-5.51%		
Last	0.4111	-17.78%	0.5135	-4.68%	0.8713	-6.79%	0.5986	-9.00%		
Random	0.4667	-6.67%	0.5312	-1.40%	0.9254	-1.00%	0.6411	-2.54%		
Ours (Step-level)	0.5222	+4.44%	<u>0.5640</u>	+4.70%	0.9387	+0.43%	<u>0.6750</u>	+2.61%		
		Qwen3-4l	B-Base &	OpenThou	ght3-100K	-				
Base	0.3667	_	0.4865	_	0.9160	_	0.5897	_		
First	0.3889	+6.05%	0.4646	-4.50%	0.8913	-2.69%	0.5816	-1.38%		
Last	0.2889	-21.20%	0.4495	-7.60%	0.8420	-8.08%	0.5268	-10.68%		
Random	0.3500	-4.56%	0.4748	-2.43%	0.8987	-1.88%	0.5745	-2.49%		
Ours (Step-level)	<u>0.4000</u>	+9.08%	<u>0.4899</u>	+0.70%	0.9113	-0.51%	<u>0.6004</u>	+1.81%		

answer after observing both the initial problem setup and the concluding consolidation region, while discarding potentially redundant intermediate reasoning.

The main results in Table 4 are obtained with <code>Qwen3-8B-Base</code> trained on OpenThought3-100K, while additional results with <code>Qwen2.5-32B-Instruct</code> trained on s1K-1.1 are provided in Table C.1. These additional results reinforce the observation from the main text that trimming 25-75% of intermediate steps yields nearly identical average performance compared to using full trajectories, indicating that <code>MidCut-Decoding</code> can reduce inference costs without harming task accuracy.

Table C.1: **Performance of MidCut-Decoding.** "Accuracy" denotes the averaged accuracy across AIME24, GPQA-D, and MATH.

Model	Full	25%	33%	50%	75%
Qwen3-8B-Base	0.6578	0.6574	0.6579	0.6581	0.6576
Qwen2.5-32B-Instruct	0.7420	0.7421	0.7420	0.7420	0.7416

D THE USE OF LLMS

LLMs were primarily used for minor language editing, including adjustments to word choices and clarity. They played no role in the research design, analysis, interpretation, or manuscript preparation, and all scientific contributions are fully our own.